

NASA Contractor Report 195039

ICASE Report No. 95-9

1N-61
W. H. H. H. H.
1.21



ICASE

DISCRETE-TIME MARKOVIAN STOCHASTIC PETRI NETS

N95-23613

Unclass

G3/61 0044074

Gianfranco Ciardo

(NASA-CR-195039) DISCRETE-TIME
MARKOVIAN STOCHASTIC PETRI NETS
Final Report (ICASE) 21 p

Contract No. NAS1-19480
February 1995

Institute for Computer Applications in Science and Engineering
NASA Langley Research Center
Hampton, VA 23681-0001



Operated by Universities Space Research Association

Discrete-time Markovian stochastic Petri nets

Gianfranco Ciardo

Department of Computer Science

College of William and Mary

Williamsburg, VA 23187-8795

ciardo@cs.wm.edu

Abstract

We revisit and extend the original definition of discrete-time stochastic Petri nets, by allowing the firing times to have a “defective discrete phase distribution”. We show that this formalism still corresponds to an underlying discrete-time Markov chain. The structure of the state for this process describes both the marking of the Petri net and the phase of the firing time for of each transition, resulting in a large state space. We then modify the well-known power method to perform a transient analysis even when the state space is infinite, subject to the condition that only a finite number of states can be reached in a finite amount of time. Since the memory requirements might still be excessive, we suggest a bounding technique based on truncation.

¹This research was partially supported by the National Aeronautics and Space Administration under NASA Contract No. NAS1-19480 while in residence at the Institute for Computer Applications in Science and Engineering (ICASE), NASA Langley Research Center, Hampton, VA 23681-0001.

1 Introduction

In the past decade, stochastic Petri nets (SPNs) have received much attention from researchers in the performance and reliability arena and have been extensively applied to the performance and reliability modeling of computer, communication, manufacturing, and aerospace systems [4, 5, 7, 10, 23]. While there is agreement on the appropriateness of SPNs as a description formalism for a large class of systems, two radically different solution approaches are commonly employed: simulation and state-space-based analysis. Simulation allows to associate general distributions to the duration of activities (SPN transitions), but it requires multiple runs to obtain meaningful statistics. This problem is particularly acute in reliability studies, where many runs might be required to obtain tight confidence intervals. With simulation, the state of the SPN is composed of the *marking*, describing the structural state of the SPN, and the *remaining firing times*, describing how long each transition in the SPN must still remain enabled before it can fire. The simulated time θ is advanced by *firing* the transition with the smallest remaining firing time.

State-space-based analysis has been mostly applied to SPNs whose underlying process is a continuous-time Markov chain (CTMC), that is, to SPNs with exponentially distributed firing times [3, 12, 25, 26]. Except for numerical truncation and roundoff, exact results are obtained, but the approach has two limitations: the number of SPN markings increases combinatorially, rendering unfeasible the solution of large models, and generally-distributed activities must be modeled using “phase-type (PH) expansion” [15]. PH distributions can approximate any distribution arbitrarily well, but it is difficult to exploit this fact in practice because the expansion exacerbates the state-space size problem.

Discrete distributions for the timing of SPNs have received less attention. This is unfortunate, since deterministic distributions (constants) are often needed to model low-level phenomena in both hardware and software, and the geometric distribution is the discrete equivalent of the exponential distribution and can approximate it arbitrarily well as the size of the step decreases. Furthermore, there is evidence supporting the use of deterministic instead of exponential distributions when modeling parallel programs [1].

If all the firing distributions are geometric with the same step, the underlying process is a discrete-time Markov chain (DTMC) [25]. Such SPNs can model synchronous behavior, as well as the main aspect of asynchronous systems: the uncertainty about the ordering of quasi-simultaneous events. A DTMC is described by a square one-step state transition probability matrix Π and an initial state probability vector $\pi^{[0]}$. The state probability vector at step n can be obtained with the iteration (*power method*): $\pi^{[n]} = \pi^{[n-1]}\Pi$. This result was extended in [11] to include immediate transitions, which fire in zero time, and geometric firing distributions with steps multiple of a basic unit step, possibly with parameter equal one, that is, constants. [29] restates these results in more detail, and uses the concept of

weight to break the ties, following [3] and, more closely, [13]. Generalized Timed Petri Nets (GTPN) have also been proposed [19], where the steps of the geometric firing times for each transition can be arbitrary, unrelated, real numbers. A DTMC can be obtained by embedding, but the analysis is restricted to steady-state behavior and the state space of the DTMC can be infinite even when the underlying untimed PN has a finite reachability set. Analogous considerations hold for D-timed PNs [30].

We generalize and formalize the results in [13] and show how, using phase-expansion, a DTMC can be obtained even if the firing time distributions are not geometric, as long as firings can occur only at some multiple of a unit step. The state can then be described by the marking plus the phase of each transition. This extends the class of SPNs that can be solved analytically, but two limitations still exist: the existence of a basic step and the size of the state space. By using a fine step, arbitrary steps can be approximated, but this increases the state space.

Approaches to solve models with a large state space have been proposed for both steady-state and transient analysis. [6] considers the reliability study of a SPN with exponentially distributed firing times, under the condition that the reachability graph is acyclic. The underlying CTMC is then acyclic as well, and a state can be discarded as soon as the transitions out of it have been explored, resulting in an algorithm offering large savings in memory and computations with respect to traditional numerical approaches. However, acyclic state spaces arise only in special cases, such as reliability models of non-repairable systems.

For transient analysis of a general CTMC, Jensen's method [21], also called uniformization [17, 27], is widely adopted. [18] outlines a dynamic implementation of the algorithm, where the state space is explored as the computation of the transient probability vector proceeds, not in advance, as normally done. This allows to obtain a transient solution even if the state space is infinite, provided that the transition rates have an upper bound.

If the CTMC contains widely different rates, the number of matrix-vector multiplications required by uniformization can be excessive. Proposals to alleviate this problem are selective randomization [24] and adaptive uniformization [28], both based on the idea of allowing different uniformization rates, according to the set of states that can be reached at each step. The latter, in addition, can be used with infinite state spaces even if the rates have no upper bound. However, the method can incur a substantial overhead, and it appears that an adaptive step is advantageous only in special cases or for short time horizons.

In Sections 2, 3, and 4 we define the underlying untimed PN model, the class of DDP distributions used for the temporization of a PN, and the resulting DDP-SPN formalism, respectively. Section 5 discusses the numerical solution of a DDP-SPN, by building and solving its underlying stochastic process, a DTMC. Section 6, examines approaches to cope with large state spaces.

2 The PN formalism

We recall the (extended) PN formalism used in [12, 14]. A PN is a tuple $(P, T, D^-, D^+, D^\circ, \succ, g, \mu^{[0]})$ where:

- P is a finite set of places, which can contain tokens. A marking $\mu \in \mathbb{N}^{|P|}$ defines the number of tokens in each place $p \in P$, indicated by μ_p (when relevant, a marking should be considered a column vector). D^- , D^+ , D° , and g are “marking-dependent”, that is, they are specified as functions of the marking.
- T is a finite set of transitions. $P \cap T = \emptyset$.
- $\forall p \in P, \forall t \in T, \forall \mu \in \mathbb{N}^{|P|}$, $D_{p,t}^-(\mu) \in \mathbb{N}$, $D_{p,t}^+(\mu) \in \mathbb{N}$, and $D_{p,t}^\circ(\mu) \in \mathbb{N}$ are the multiplicities of the input arc from p to t , the output arc from t to p , and the inhibitor arc from p to t , when the marking is μ , respectively.
- $\succ \subseteq T \times T$ is an acyclic (pre-selection) priority relation.
- $\forall t \in T, \forall \mu \in \mathbb{N}^{|P|}$, $g_t(\mu) \in \{0, 1\}$ is the guard for t in marking μ .
- $\mu^{[0]} \in \mathbb{N}^{|P|}$ is the initial marking.

Places and transitions are drawn as circles and rectangles, respectively. The number of tokens in a place is written inside the place itself (default is zero). Input and output arcs have an arrowhead on their destination, inhibitor arcs have a small circle. The multiplicity is written on the arc (default is the constant 1); a missing arc indicates that the multiplicity is the constant 0. The default value for guards is the constant 1.

A transition $t \in T$ is enabled in marking μ iff all the following conditions hold:

1. $g_t(\mu) = 1$.
2. $\forall p \in P, D_{p,t}^-(\mu) \leq \mu_p$.
3. $\forall p \in P, D_{p,t}^\circ(\mu) > \mu_p$ or $D_{p,t}^\circ(\mu) = 0$.
4. $\forall u \in T, u \not\succ t$ or u is not enabled in μ (well defined because \succ is acyclic).

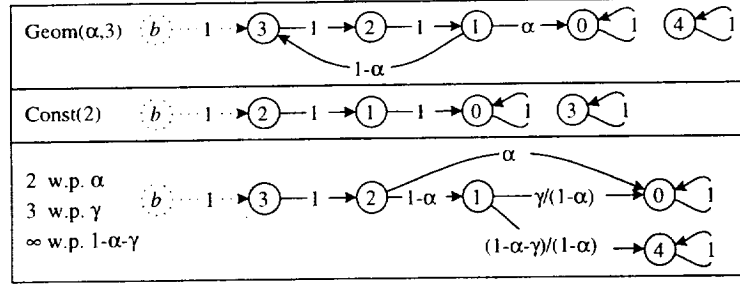


Figure 1: Examples of DDP distributions.

Let $\mathcal{E}(\mu)$ be the set of transitions enabled in marking μ . A transition $t \in \mathcal{E}(\mu)$ can fire, causing a change to marking $\mathcal{M}(t, \mu)$, obtained from μ by subtracting the “input bag” $D_{\bullet, t}^-(\mu)$ and adding the “output bag” $D_{\bullet, t}^+(\mu)$ to it: $\mathcal{M}(t, \mu) = \mu - D_{\bullet, t}^-(\mu) + D_{\bullet, t}^+(\mu) = \mu + D_{\bullet, t}(\mu)$, where $D = D^+ - D^-$ is the incidence matrix. \mathcal{M} can be extended to its reflexive and transitive closure by considering the marking reached from μ after firing a sequence of transitions. The *reachability set* is given by $\mathcal{R} = \{\mu : \exists \sigma \in T^* \wedge \mu = \mathcal{M}(\sigma, \mu^{[0]})\}$, where T^* indicates the set of transition sequences.

3 Discrete phase distributions

We now define the class \mathcal{D} of (possibly defective) discrete phase (DDP) distributions, which will be used to specify the duration of a firing time in a SPN. A random variable X is said to have a DDP distribution, $X \sim \mathcal{D}$, iff there exists an absorbing DTMC $\{A^{[k]} : k \in \mathbb{N}\}$ with finite state space $\mathcal{A} = \{0, 1, \dots, n\}$ and initial probability distribution given by $[\Pr\{A^{[0]} = i\}, i \in \mathcal{A}]$, such that states $\mathcal{A} \setminus \{0, n\}$ are transient and states $\{0, n\}$ are absorbing, and X is the time to reach state 0: $X = \min\{k \geq 0 : A^{[k]} = 0\}$. If $\Pr\{A^{[0]} = 0\} > 0$, the distribution has a mass at the origin. If $\Pr\{A^{[0]} = i\} > 0$ and state i can reach state n , the distribution is (strictly) defective.

\mathcal{D} is the smallest class containing the distributions $\text{Const}(0)$, $\text{Const}(1)$, and $\text{Const}(\infty)$ and closed under:

- Finite convolution: if $X_1 \sim \mathcal{D}$ and $X_2 \sim \mathcal{D}$, then $X = X_1 + X_2 \sim \mathcal{D}$.
- Finite weighted sum: if $X_1 \sim \mathcal{D}$, $X_2 \sim \mathcal{D}$ and $B \in \{0, 1\}$ is a Bernoulli random variable, then $X = BX_1 + (1 - B)X_2 \sim \mathcal{D}$.
- Infinite geometric sum: if $\{X_k \sim \mathcal{D} : k \in \mathbb{N}^+\}$ is a family of iid's and N is a geometric random variable, then $X = \sum_{1 \leq k \leq N} X_k \sim \mathcal{D}$.




Figure 2: Equivalent DTMC representations

The geometric and modified geometric distributions with arbitrary positive integer step, $\text{Geom}(\alpha, \omega)$ and $\text{ModGeom}(\alpha, \omega)$, $0 \leq \alpha \leq 1$, $\omega \in \mathbb{N}^+$, the constant non-negative integer distribution, $\text{Const}(\omega)$, $\omega \in \mathbb{N}$, and any discrete distribution with finite non-negative integer support are special cases of DDP distributions. An example of a random variable with non-negative integer support which does not have a DDP distribution is N^2 , where $N \sim \text{Geom}(\alpha)$.

Fig. 1 shows examples of DDP distributions. The “initial state” b , for begin, has zero sojourn time and is introduced to represent graphically the initial probability distribution. We use this representation since it allows to estimate the “complexity” of a DTMC by counting the number of nodes and arcs in its graph. For simplicity, the last state, e.g., 4 for $\text{Geom}(\alpha, 3)$ and 3 for $\text{Const}(2)$, can be omitted if it is not reachable from b (if the distribution is actually not defective). Unfortunately, the DTMC corresponding to a given DDP distribution might not be unique, even if the number of states is fixed. For example, the time X to reach state 0 for the DTMCs in Fig. 2, both with five nodes and seven arcs, has distribution $\text{Unif}(0, 3)$, that is, $\Pr\{X = i\} = 1/4$, for $i \in \{0, 1, 2, 3\}$.

4 The DDP-SPN formalism

SPNs are obtained when the time that must elapse between the instant a transition becomes enabled and the instant it can fire, or firing time, is a random variable. By restricting the firing times distributions to \mathcal{D} , we obtain the DDP-SPNs, corresponding to a stochastic process where the state has the form $s = (\mu, \phi) \in \mathbb{N}^{|P|} \times \mathbb{N}^{|T|}$. The structural component μ is simply the current marking. The timing component ϕ describes the current “phases”, the state for the DTMC chosen to encode the DDP distribution associated with the firing time of each transition. The firing time of a transition t elapses when its phase ϕ_t reaches 0. Formally, a DDP-SPN is a tuple $(P, T, D^-, D^+, D^\circ, \succ, g, \mu^{[0]}, \Phi, G, F, \phi^{[0]}, \succ, w)$ where:

- $(P, T, D^-, D^+, D^\circ, \succ, g, \mu^{[0]})$ define a PN.
- $\forall t \in T, \forall \mu \in \mathcal{R}, \Phi_t(\mu) \subset \mathbb{N}$ is the finite set of possible phases in which transition t can be when the marking is μ .

- $\forall \mu \in \mathcal{R}, \forall t \in \mathcal{E}(\mu), \forall i, j \in \Phi_t(\mu), G_t(\mu, i, j)$ is the probability that the phase of t changes from i to j at the end of one step, when t is enabled in marking μ . Hence, $\sum_{j \in \Phi_t(\mu)} G_t(\mu, i, j) = 1$.
- $\forall \mu \in \mathcal{R}, \forall u \in \mathcal{E}(\mu), \forall t \in T, \forall i \in \Phi_t(\mu), \forall j \in \Phi_t(\mathcal{M}(u, \mu)), F_{u,t}(\mu, i, j)$ is the probability that the phase of t changes from i to j when u fires in marking μ . Hence, $\sum_{j \in \Phi_t(\mathcal{M}(u, \mu))} F_{u,t}(\mu, i, j) = 1$.
- $\forall t \in T, \phi_t^{[0]} \in \Phi_t(\mu^{[0]})$ is the phase of t at time 0.
- $\succ \subseteq T \times T$ is an acyclic (post-selection) priority relation.
- $\forall \mu \in \mathcal{R}, \forall S \subseteq \mathcal{E}(\mu), \forall t \in S, w_{t|S}(\mu) \in \mathbb{R}^+$ is the firing weight for t when S is the set of candidates to fire in marking μ .

A transition $t \in T$ is said to be a *candidate* (to fire) in state $s = (\mu, \phi)$ iff all the following conditions hold:

1. $t \in \mathcal{E}(\mu)$.
2. $\phi_t = 0$.
3. $\forall u \in T, u \not\succ t$ or u is not a candidate in s (remember that \succ is acyclic).

Let $\mathcal{C}(s)$ be the set of candidates in state s . $G_t(\mu, \bullet, \bullet)$ is the one-step transition probability matrix of the DTMC $\{\phi_t^{[k]} : k \in \mathbb{N}\}$, with state space $\Phi_t(\mu)$, corresponding to the DDP-distributed firing time for transition t in marking μ in isolation, that is, assuming that no other transition firing affects the firing time of t . However, if another transition u fires before t , leading to marking μ' , the phase ϕ_t of t will change according to the distribution $F_{u,t}(\mu, \phi_t, \bullet)$. Furthermore, after the firing of u , the phase of t will evolve according to $G_t(\mu', \bullet, \bullet)$, which might differ from $G_t(\mu, \bullet, \bullet)$, it can even have a different state space, $\Phi_t(\mu')$ instead of $\Phi_t(\mu)$.

We stress that pre-selection and post-selection have a different semantic. Only in the case of immediate transitions the two become equivalent. Assume that only t and u satisfy the input, inhibitor, and guard conditions in μ . We have three options, resulting in three different behaviors:

- Specify a pre-selection priority between them, for example $t \succ u$, so that u will not be enabled when t is. This means that the phase ϕ_t of t evolves according to $G_t(\mu, \bullet, \bullet)$, while ϕ_u does not. The same effect would be achieved using a guard $g_u(\mu) = 0$.

- Specify no pre-selection priority, but a post-selection priority between them, for example $t \succ u$. This means that the phases of both t and u evolve in μ . The first one to reach phase 0 will fire but, in case of a tie, t will be chosen. However, if $\phi_u = 0$ when t fires and if $F_{t,u}(\mu, 0, 0) = 1$, u might be a candidate in the new marking, and fire immediately after t .
- Specify neither a pre-selection nor a post-selection priority between them. Then, as in the previous case, t and u are in a race to reach phase 0, but a tie is now resolved by a probabilistic choice according to the weights: $\hat{w}_{t|\{t,u\}}(\mu)$ and $\hat{w}_{u|\{t,u\}}(\mu)$, respectively, where \hat{w} is a normalization of w to ensure that the weights of the candidates in a marking sum to one.

Let $(\mu^{[n]}, \phi^{[n]})$ be the state of the DDP-SPN at step n . Then, the process $\{(\mu^{[n]}, \phi^{[n]}) : n \in \mathbb{N}\}$ is a DTMC with state space $\mathcal{S} \subseteq \mathbb{N}^{|P|} \times \mathbb{N}^{|T|}$. Its one-step transition probability matrix Π is determined by considering the possibility of simultaneous firings. Consider a state $s = (\mu, \phi)$. If $\mathcal{C}(s) \neq \emptyset$, one of the candidates will fire immediately, and the sojourn time in s is zero. Otherwise, the sojourn time in s is one. Following GSPN [3] terminology, we call s a vanishing or tangible state, respectively. Hence, s is tangible iff $\phi > 0$.

Let $S_{s,s'}$ be the set of possible event sequences events leading from a tangible state $s = (\mu, \phi)$ to a tangible state $s' = (\mu', \phi')$ in one time step:

$$S_{s,s'} = \left\{ \sigma = (\mu^{(0)}, \phi^{(0)}, t^{(0)}, \mu^{(1)}, \phi^{(1)}, t^{(1)}, \dots, \mu^{(n-1)}, \phi^{(n-1)}, t^{(n-1)}, \mu^{(n)}, \phi^{(n)}) : \right. \\ \left. n \geq 0, \mu^{(0)} = \mu, \mu^{(n)} = \mu', \phi^{(n)} = \phi', \right. \\ \left. \forall t \in \mathcal{E}(\mu), G_t(\mu, \phi_t, \phi_t^{(0)}) > 0, \right. \quad (1)$$

$$\left. \forall i, 0 \leq i < n, t^{(i)} \in \mathcal{C}(\mu^{(i)}, \phi^{(i)}), \mu^{(i+1)} = \mathcal{M}(t^{(i)}, \mu^{(i)}), \right. \quad (2)$$

$$\left. \forall t \in T, F_{t^{(i)}, t}(\mu^{(i)}, \phi_t^{(i)}, \phi_t^{(i+1)}) > 0 \right\}. \quad (3)$$

(1) considers the one-step evolution of the phases for the enabled transitions in isolation, while (2) and (3) consider the sequentialized firing in zero time of zero or more transitions at the end of the one-step period. Hence, $(\mu^{(i)}, \phi^{(i)})$ is a vanishing state, for $0 \leq i < n$.

The value of the nonzero entries of Π is obtained by summing the probability of all possible sequences leading from s to s' :

$$\Pi_{s,s'} = \sum_{\sigma \in S_{s,s'}} \left(\prod_{t \in \mathcal{E}(\mu)} G_t(\mu, \phi_t, \phi_t^{(0)}) \right) \\ \cdot \left(\prod_{i=0}^{n-1} \hat{w}_{t^{(i)}|\mathcal{C}(\mu^{(i)}, \phi^{(i)})}(\mu^{(i)}) \left(\prod_{t \in T} F_{t^{(i)}, t}(\mu^{(i)}, \phi_t^{(i)}, \phi_t^{(i+1)}) \right) \right)$$

In a practical implementation, Π is computed one row at a time. The complexity of computing row s of Π can be substantial, depending on the length and number of sequences

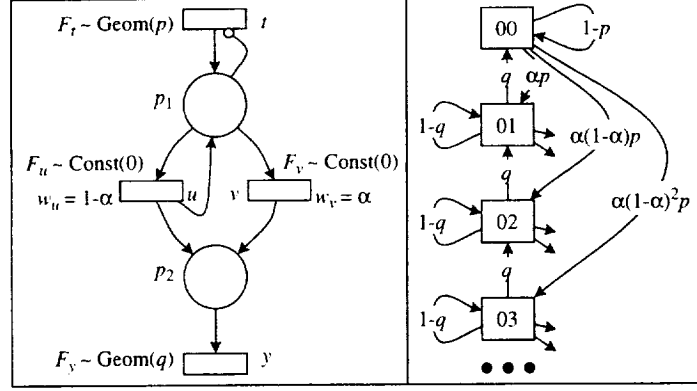


Figure 3: $(0, 0)$ can reach an infinite number of markings in one time step.

in $\bigcup_{s'} S_{s,s'}$. If $\bigcup_{s'} S_{s,s'}$ is infinite, special actions must be taken. This can happen for two reasons:

- \mathcal{R} is itself infinite, and state s can reach an infinite number of states in a single step. Consider, for example, a single queue with batch arrivals of size $N > 0$, where $N \sim \text{Geom}(\alpha)$, as in Fig. 3. Following the firing of t , a geometrically distributed number of tokens will be placed in p_2 : when the token is finally removed from p_1 (by the firing of v), p_2 contains N tokens with probability $\alpha(1-\alpha)^{N-1}$. This represents a batch arrival of size N at the server modeled by place p_2 and transition y . Unfortunately, finiteness of \mathcal{R} is an undecidable question for the class of Petri nets we defined, since transition priorities alone make them Turing equivalent [2].
- $S_{s,s'}$ can be infinite for a particular s' . If \mathcal{R} is finite, this requires the presence of arbitrarily long paths over a finite set of vanishing states, just as for a “vanishing loop” in a GSPN [11]. In a practical implementation, these cycles can be detected and managed appropriately.

The size of the DTMC underlying a DDP-SPN is affected by the choice of the representation for the DDP distributions involved. Consider, for example, the DDP-SPN in Fig. 4(a), and assume that transitions t_1 , t_2 , and t_3 have firing time distributions $\text{Const}(1)$, $\text{Unif}(0, 3)$, and $\text{Const}(2)$, respectively. The corresponding DTMCs obtained using the two representations of Fig. 2 for $\text{Unif}(0, 3)$ are shown in Fig. 4(b) and 4(c), respectively. The number of states is ten in the first case, seven in the second (the value of ϕ_t is specified as “•” whenever t is not enabled and either it cannot become enabled again or its phase is going to be reset upon becoming enabled). The difference between the size of the two DTMCs is due to a *lumping* [22] of the states, and it would be even greater if t_3 had a more complex distribution. By postponing the probabilistic decision as much as possible, the second DTMC lumps




Figure 4: The effect of equivalent $\text{Unif}(0, 3)$ representations.

states $(011, \bullet 12)$, $(011, \bullet 22)$, and $(011, \bullet 32)$ of the first DTMC into a single one, $(011, \bullet 32)$, and states $(011, \bullet 11)$ and $(011, \bullet 21)$ into $(011, \bullet 21)$.

5 Analysis of DDP-SPNs


When using a SPN to model a system, a reward rate ρ_μ is associated to each marking μ . Starting from $\{(\mu^{[n]}, \phi^{[n]}) : n \in \mathbb{N}\}$, it is then possible to define two continuous-parameter processes: $\{y(\theta), \theta \geq 0\}$, describing the instantaneous reward rate at time θ : $y(\theta) = \rho_{\mu(\theta)}$, where $\mu(\theta) = \mu^{\lfloor \max\{n \leq \theta\} \rfloor}$, and $\{Y(\theta), \theta \geq 0\}$, describing the reward accumulated up to time θ , $Y(\theta) = \int_0^\theta \rho_{\mu(\tau)} d\tau$.

We consider the computation of the expected value of $y(\theta_F)$ and $Y(\theta_F)$ for finite values of θ_F . Let $\pi^{[n]} = [\pi_s^{[n]}] = [\Pr\{s^{[n]} = s\}]$ be the state probability vector at time n . Once the state-space \mathcal{S} corresponding to the initial state $(\mu^{[0]}, \phi^{[0]})$ has been generated, any initial probability vector over \mathcal{S} can be used for the initial probability vector $\pi^{[0]}$, there is no requirement to use a vector having a one in position $(\mu^{[0]}, \phi^{[0]})$ and a zero elsewhere. From $\pi^{[0]}$, we can obtain $\pi^{[n]}$ iteratively, performing n matrix-vector multiplications:

$$\pi^{[n]} = \pi^{[n-1]} \Pi \quad (4)$$

Since the DTMC can change state only at integer times, $\pi(\theta) = \pi^{[n]}$ for $\theta \in [n, n+1)$. Practical implementations assume that the state space is finite and that the transition probability matrix Π is computed before starting the iterations. The following shows the pseudo-code to compute $E[y(\theta_F)]$ and $E[Y(\theta_F)]$ with the “power method”:

1. “compute \mathcal{S} , Π , and $\pi^{[0]}$ ”;
2. $Y \leftarrow 0$; $\pi \leftarrow \pi^{[0]}$;
3. for $n = 1$ to $\lfloor \theta_F \rfloor$ do



REPORT DOCUMENTATION PAGE		Form Approved OMB No. 0704-0188
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.		
1. AGENCY USE ONLY(Leave blank)	2. REPORT DATE February 1995	3. REPORT TYPE AND DATES COVERED Contractor Report
4. TITLE AND SUBTITLE DISCRETE-TIME MARKOVIAN STOCHASTIC PETRI NETS		5. FUNDING NUMBERS C NAS1-19480 WU 505-90-52-01
6. AUTHOR(S) Gianfranco Ciardo		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Computer Applications in Science and Engineering Mail Stop 132C, NASA Langley Research Center Hampton, VA 23681-0001		8. PERFORMING ORGANIZATION REPORT NUMBER ICASE Report No. 95-9
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Langley Research Center Hampton, VA 23681-0001		10. SPONSORING/MONITORING AGENCY REPORT NUMBER NASA CR-195039 ICASE Report No. 95-9
11. SUPPLEMENTARY NOTES Langley Technical Monitor: Dennis M. Bushnell Final Report To appear in Numerical Solution of Markov Chains '95		
12a. DISTRIBUTION/AVAILABILITY STATEMENT Unclassified-Unlimited Subject Category 60, 61		12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 words) We revisit and extend the original definition of discrete-time stochastic Petri nets, by allowing the firing times to have a "defective discrete phase distribution". We show that this formalism still corresponds to an underlying discrete-time Markov chain. The structure of the state for this process describes both the marking of the Petri net and the phase of the firing time for of each transition, resulting in a large state space. We then modify the well-known power method to perform a transient analysis even when the state space is infinite, subject to the condition that only a finite number of states can be reached in a finite amount of time. Since the memory requirements might still be excessive, we suggest a bounding technique based on truncation.		
14. SUBJECT TERMS Markov Chains		15. NUMBER OF PAGES 20